arXiv:1408.1789v3 [cs.CG] 6 Dec 2015

Dimension reduction techniques for ℓ_p $(1 \leq p \leq 2)$, with applications

Yair Bartal * Lee-Ad Gottlieb †

December 8, 2015

Abstract

For Euclidean space (ℓ_2), there exists the powerful dimension reduction transform of Johnson and Lindenstrauss [JL84], with a host of known applications. Here, we consider the problem of dimension reduction for all ℓ_p spaces $1 \leq p \leq 2$. Although strong lower bounds are known for dimension reduction in ℓ_1, Ostrovsky and Rabani [OR02] successfully circumvented these by presenting an ℓ_1 embedding that maintains fidelity in only a bounded distance range, with applications to clustering and nearest neighbor search. However, their embedding techniques are specific to ℓ_1 and do not naturally extend to other norms.

In this paper, we apply a range of advanced techniques and produce bounded range dimension reduction embeddings for all of $1 \leq p \leq 2$, thereby demonstrating that the approach initiated by Ostrovsky and Rabani for ℓ_1 can be extended to a much more general framework. We also obtain improved bounds in terms of the intrinsic dimensionality. As a result we achieve improved bounds for proximity problems including snowflake embeddings and clustering.

*Hebrew University. Work supported in part by an Israel Science Foundation grant #1609/11. Email: yair@cs.huji.ac.il

†Ariel University. Email: leead@ariel.ac.il

1 Introduction

Dimension reduction for normed space is a fundamental tool for algorithms and related fields. A much celebrated result for dimension reduction is the well-known l_2 flattening lemma of Johnson and Lindenstrauss [JL84]: For every n-point subset of l_2 and every $0 < \varepsilon < 1$, there is a mapping into l_2^k that preserves all interpoint distances in the set within factor $1 + \varepsilon$, with target dimension $k = O(\varepsilon^{-2} \log n)$. The dimension reducing guarantee of the Johnson-Lindenstrauss (JL) transform is remarkably strong, and has the potential to make algorithms with a steep dependence on the dimension tractable. It can be implemented as a simple linear transform, and has proven to be a very popular tool in practice, even spawning a stream of literature devoted to its analysis, implementation and extensions (see for example [Alo03, AC06, AL11, Ach01, DKS10, VW11]).

Given the utility and impact of the JL transform, it is natural to ask whether these strong dimension reduction guarantees may hold for other ℓ_p spaces as well (see [Ind01] for further motivation). This is a fundamental open problem in embeddings, and has attracted significant research. Yet the dimension reduction bounds known for ℓ_p norms other than ℓ_2 are much weaker than those given by the JL transform, and it is in fact known that a linear dimension reduction mapping in the style the JL transform is quite unique to the ℓ_2-norm [JN09]. Further, strong lower bounds on dimension reduction are known for ℓ_1 [BC05, LN04, ACNN11] and for ℓ_∞ [Mat96], and it is a plausible conjecture that ℓ_2 is the only ℓ_p space which admits the strong distortion and dimension properties of the JL transform.

Ostrovsky and Rabani [KOR98, OR00, OR02] successfully circumvented the negative results for ℓ_1, and presented a dimension reduction type embedding for the Hamming cube – and by extension, all of l_1 – which preserves fidelity only in a bounded range of distances. They further demonstrated that their embedding finds use in algorithms for nearest neighbor search (NNS) and clustering, as these can be reduced to subproblems where all relevant distance are found in a bounded range. In fact, a number of other proximity problems can also be reduced to bounded range subproblems, including the construction of distance oracles and labels [HM06, BGK$^+$11], snowflake embeddings [GK11, BRS11], and ℓ_p-difference of a pair of data streams. Hence, we view a bounded range embedding as a framework for the solution of multiple important problems. Note also that for spaces with fixed aspect ratio (a fixed ratio between the largest and smallest distances in the set – a common scenario in the literature, see [IN07]), a bounded range mapping is in effect a complete dimension reduction embedding.

The dimension reduction embedding of Ostrovsky and Rabani is very specific to ℓ_1, and does not naturally extend to other norms. The central contribution of the paper is to bring advanced techniques to bear on this problem, thereby extending this framework to all $1 \leq p \leq 2$.

Our contribution: We first present a basic embedding in Section 3, which shows that we can reduce dimension while realizing a certain distance transform with low distortion. Using this result, we are able to derive two separate dimension-reducing embeddings:

- Range embedding. In Theorem 4.1, we present an embedding which preserves distances in a given range with $(1+\varepsilon)$ distortion. The target dimension is $O(\log n)$, with dependence on the range parameter and ε. This generalizes the approach of Ostrovsky and Rabani [OR00, OR02] to all $1 \leq p \leq 2$. This embedding can be applied in the streaming setting as well, and it can also be modified to achieve target dimension polynomial in the *doubling dimension* of the set

(Lemma 4.2).[1]

- Snowflake embedding. An α-snowflake embedding is one in which each inter-point distance t in the origin space is replaced by distance t^α in the embedded space (for $0 < \alpha < 1$). It was observed in [GK11, BRS11] that the snowflake of a finite metric space in l_2 may be embedded in dimension which is close to the intrinsic dimension of the space (measured by its doubling dimension), and this may be independent of n. In [GK11] the case of l_1 was considered as well, however the resulting dimension had doubly exponential dependence on the doubling dimension. We demonstrate that the basic embedding can be used to build a snowflake for ℓ_p for all $1 \le p \le 2$ with dimension polynomial in the doubling dimension; this is found in Lemma 4.4. For ℓ_1 this provides a *doubly exponential* improvement over the previously known dimension bound [GK11], while generalizing the $p \in \{1, 2\}$ results of [GK11, BRS11] to all $1 \le p \le 2$.

Application to clustering. To highlight the utility of our embeddings, we demonstrate (in Section 5) their applicability in deriving better runtime bounds for clustering problems:

We first consider the k-center problem, and show that our snowflake embedding can be used to provide an efficient algorithm. For ℓ_p-spaces of low doubling dimension and fixed p, $1 \le p \le 2$, we apply our range embedding in conjunction with the clustering algorithm of Agarwal and Procopiuc [AP02], and obtain a $(1+\varepsilon)$-approximation to the k-center problem in time $O(n(2^{\tilde{O}(\mathrm{ddim}(S))} + \log k)) + (k \cdot \varepsilon^{-\mathrm{ddim}(S)/\varepsilon^2})^{k^{1-(\varepsilon/\mathrm{ddim}(S))^{O(1)}}}$.

We then consider the min-sum clustering problem, and show that our snowflake embedding can be used to provide an efficient algorithm for this problem. For ℓ_p-spaces of low doubling dimension, we apply our snowflake embedding in conjunction with the clustering algorithm of Schulman [Sch00], and obtain a $(1+\varepsilon)$-approximation to the k-center problem in randomized time $n^{O(1)} + 2^{2^{(O(d'))}}(\varepsilon \log n)^{O(d')}$ where $d' = (\mathrm{ddim}/\varepsilon)^{O(1)}$.

Related work. For results on dimension reduction for ℓ_p spaces, see [Sch11] for $p < 2$, and also [Sch87, Tal90] for $p = 1$, [Bal90, Tal95] for $1 < p < 2$, and [Mat96] for $p = \infty$. Other related notions of dimension reduction have been suggested in the literature. Indyk [Ind06] devised an analogue to the JL-Lemma which uses p-stable distributions to produce estimates of interpoint distances; strictly speaking, this is not an embedding into ℓ_p (e.g. it uses median over the coordinates). Motivated by the nearest neighbor search problem, Indyk and Naor [IN07] proposed a weaker form of dimension reduction, and showed that every doubling subset $S \subset \ell_2$ admits this type of dimension reduction into ℓ_2 of dimension $O(\mathrm{ddim}(S))$. Roughly speaking, this notion is weaker in that distances in the target space are allowed to err in one direction (be too large) for all but one pair of points. Similarly, dimension reduction into ordinal embeddings (where only relative

[1] Our range embedding has the additional property that it can be used to embed ℓ_p^m into $\ell_q^{O(\log n)}$ (that is, m-dimensional ℓ_p into $O(\log n)$-dimensional ℓ_q, for $1 \le q < p$) with $(1+\varepsilon)$-distortion in time $O(m \log n)$, with dependence on the range parameter and ε. This is a fast version of the embedding of [JS82], a common tool for embedding ℓ_p into more malleable spaces such as l_1. (The embedding of [JS82] is particularly useful for nearest neighbor search, see [KOR98, HIM12].) Note that [JS82] features a large overhead cost $O(m^2)$, and since m can be as large as $\Theta(n)$, this overhead can be the most expensive step in algorithms for ℓ_p.

We also note that the embedding of Theorem 4.1 can be used to produce efficient algorithms for approximate nearest neighbor search and ℓ_p difference, although other efficient techniques have already been developed for these specific problems (see [Kle97, IM98, KOR98, And09, DIIM04, Pan06, Ngu13] for NNS, and [FKSV03, FS01, AMS96, AGMS99, IW05, IW05, Ind06, Li08] for ℓ_p difference).

distance is approximately preserved) was considered in [ABD+08, BDH+08]. Bartal, Recht and Schulman [BRS11] developed a variant of the JL-Lemma that is local – it preserves the distance between every point and the $\hat{k}$ points closest to it. Assuming $S \subset \ell_2$ satisfies a certain growth rate condition, they achieve, for any desired $\hat{k}$ and $\varepsilon > 0$, an embedding of this type with distortion $1+\varepsilon$ and dimension $O(\varepsilon^{-2} \log \hat{k})$.

2 Preliminaries

Embeddings and metric transforms. Following [BES06], we define an *oblivious* embedding to be an embedding which can be computed for any point of a database set X or query set Y, without knowledge of any other point in X or Y. (This differs slightly from the definition put forth by Indyk and Naor [IN07].) Familiar oblivious embeddings include standard implementations of the JL-Lemma for l_2, the dimension reduction mapping of Ostrovsky and Rabani [OR02] for the Hamming cube, and the embedding of Johnson and Schechtman [JS82] for ℓ_p, $p \leq 2$. A *transform* is a function mapping from the positive reals to the positive reals, and a *metric transform* maps a metric distance function to another metric distance function on the same set of points. An embedding is *transform preserving* with respect to a transform if it achieves the distances defined by that transform.

Range Embedding. Let $(X, d_X), (Y, d_Y)$ be metric spaces. For distance scales $0 \leq a \leq b \leq \infty$, an $[a,b]$-embedding of X into Y with distortion D is a mapping $f : X \to Y$ such that for all $x, y \in X$ such that $d_X(x,y) \in [a,b]$: $1 \leq c \cdot \frac{d_Y(f(x),f(y))}{d_X(x,y)} \leq D$. (Here, c is any scaling constant.) Then f is a *range embedding* with range $[a,b]$. If f has the additional property that for all x, y such that $d(x,y) < a$ we have $c \cdot \frac{d_Y(f(x),f(y))}{a} \leq D$, then we say that f is range preserving from below. Similarly, if for all x, y such that $d(x,y) > b$ we have $c \cdot \frac{d_Y(f(x),f(y))}{b} \geq 1$, then we say that f is range preserving from above. And if f is range preserving from above and below, then we say that f is a *strong* range embedding.

Let $R > 1$ be a parameter. We say that X admits an R-range embedding into Y with distortion D if for every $u > 0$ there exists an $[u, uR]$ embedding of X into Y with distortion D. As above, an R-range embedding may be range preserving from above or below. We will usually take $u = 1$.

Nets and hierarchies. Given a metric space S, $S' \subset S$ is a γ-net of S if the minimum interpoint distance in S' is at least γ, while the distance from every point of S to its nearest neighbor in S' is less than γ. Let S have minimum inter-point distance 1. A *hierarchy* is a series of $\lceil \log \Delta \rceil$ nets (Δ being the aspect ratio of S), where each net S_i is a 2^i-net of the previous net S_{i-1}. The first (or bottom) net is $S_0 = S$, and the last (or top) net S_t contains a single point called the *root*. For two points $u \in S_i$ and $v \in S_{i-1}$, if $d(u,v) < 2^i$ then we say that u *covers* v, and this definition allows v to have multiple covering points in S_i. The closest covering point of v is its *parent*. The distance from a point in S_i to its ancestor in S_j is at most $\sum_{k=i+1}^{j} 2^k = 2 \cdot (2^j - 2^{i+1}) < 2 \cdot 2^j$.

Given S, a hierarchy for S can be built in time $2^{O(\mathrm{ddim}(S))} n$, and this term bounds the size of the hierarchy as well [HM06, CG06]. The height of the hierarchy is $O(\min\{n, \log \Delta\})$.

Doubling dimension. For a metric $(\mathcal{X}, \rho)$, let $\lambda > 0$ be the smallest value such that every ball in $\mathcal{X}$ can be covered by λ balls of half the radius. The *doubling dimension* of $\mathcal{X}$ is $\mathrm{ddim}(\mathcal{X}) = \log_2 \lambda$.

Note that $\mathrm{ddim}(\mathcal{X}) \leq \log n$. The following packing property can be shown (see, for example [KL04]): Suppose that $S \subset \mathcal{X}$ has a minimum interpoint distance of at least α. Then $|S| \leq \left(\frac{2\,\mathrm{diam}(S)}{\alpha}\right)^{\mathrm{ddim}(\mathcal{X})}$.

Probabilistic partitions. Probabilistic partitions are a common tool used in embeddings. Let (X, d) be a finite metric space. A partition P of X is a collection of non-empty pairwise disjoint clusters $P = \{C_1, C_2, \ldots, C_t\}$ such that $X = \cup_j C_j$. For $x \in X$ we denote by $P(x)$ the cluster containing x. We will need the following decomposition lemma due to Gupta, Krauthgamer and Lee [GKL03], Abraham, Bartal and Neiman [ABN08], and Chan, Gupta and Talwar [CGT10].[2] Let ball $B(x, r) = \{y | \, \|x - y\| \leq r\}$.

Theorem 2.1 (Padded Decomposition of doubling metrics [GKL03, ABN08, CGT10])**.** *There exists a constant $c_0 > 1$, such that for every metric space (X, d), every $\varepsilon \in (0, 1)$, and every $\delta > 0$, there is a multi-set $\mathcal{D} = [P_1, \ldots, P_m]$ of partitions of X, with $m \leq c_0 \varepsilon^{-1} \dim(X) \log \dim(X)$, such that*

1. *Bounded radius:* $\mathrm{diam}(C) \leq \delta$ *for all clusters* $C \in \bigcup_{i=1}^{m} P_i$.
2. *Ball padding: If P is chosen uniformly from $\mathcal{D}$, then for all $x \in X$,*

$$\Pr_{P \in \mathcal{D}}[B(x, \tfrac{\delta}{c_0 \dim(X)}) \subseteq P(x)] \geq 1 - \varepsilon.$$

Stable distributions. The density of a symmetric p-stable random variable ($0 < p \leq 2$) is $h(x) = \frac{1}{\pi} \int_0^\infty \cos(tx) e^{-t^p} dt$ [Zol86, Section 2.2]. The density function $h(x)$ is unimodal [SY78, Yam78, Hal84] and bell-shaped [Gaw84]. It is well known that $\frac{c_p}{1+x^{p+1}} \leq h(x) \leq \frac{c'_p}{1+x^{p+1}}$ for constants c_p, c'_p that depend only on p [Nol12] (and we may ignore the dependence on p for the purposes of this paper). Using this approximation for $h(x)$, it is easy to see that for $0 < q < p$ and p-stable random variable g, $\mathbb{E}[g^q] = \int_0^\infty x^q h(x) dx \approx \int_0^1 x^q dx + \int_1^\infty x^{q-(p+1)} dx \approx \frac{1}{p-q}$. (Here we use the notation $\approx$ in the same sense as $\Theta(\cdot)$.) Also, for $b > 1$, $\int_0^b x^p h(x) dx \approx \int_0^1 x^p dx + \int_1^b \frac{1}{x} dx \approx 1 + \ln b$. Let $g_1, g_2, \ldots, g_m \in G$ be a set of i.i.d. symmetric p-stable random variables, and let v be a real m-length vector. A central property of p-stable random variables is that $\sum_{j=1}^m g_j v_j$ is distributed as $g(\sum_{j=1}^m v_j^p)^{1/p} = g\|v\|_p$, for all $g \in G$. When all $g_i \in G$ are normalized as $\mathbb{E}[|g_i|^q] = 1$, we have that $\mathbb{E}\left[\left|\sum_{j=1}^m g_j v_j\right|^q\right] = \mathbb{E}[|g|^q]\|v\|_p^q = \|v\|_p^q$.

3 Basic transform-preserving embedding

In Theorem 3.2 below, we present an embedding which realizes a certain distance transform with low distortion, while also reducing dimension to $O(\log n)$ (with dependence on the range and the desired distortion). We then demonstrate that the transform itself has several desirable properties. In the next section, we will use this transform-preserving embedding to obtain a range embedding and a snowflake embedding.

[2] [GKL03] provided slightly different quantitative bounds than in Theorem 2.1. The two enumerated properties follow, for example, from Lemma 2.7 in [ABN08], and the bound on support-size m follows by an application of the Lovász Local Lemma sketched therein.

3.1 Embedding

We first present a randomized embedding into a single coordinate, and show that it is transform-preserving in expectation only (Lemma 3.1). We then show that a concatenation of many single-coordinate embeddings yields a single embedding which (approximately) preserves the bounded transform with high probability, and this gives Theorem 3.2.

The following single-coordinate embedding is inspired by the Nash device of Bartal *et al.* [BRS11] (see also [RR07]), and is related to the spherical threshold function of Mendel and Naor [MN04, Lemma 5.9]. Broadly speaking, our embedding uses the sine function as a dampening tool, which serves to mitigate undesirable properties of p-stable distributions (i.e., their heavy tails). Our central contribution is to give a tight analysis of the transform preserved by our embedding.[3]

Let $S \subset \ell_p$ be a set of m-dimensional vectors for $1 \leq p \leq 2$. Let $g_j \in G$ be k i.i.d. symmetric p-stable random variables, and fix parameters s (the *threshold*) and $0 \leq \phi \leq 2\pi$. Further define the fixed constant $P_q = \mathbb{E}[|\cos\theta|^q] = \frac{1}{2\pi}\int_0^{2\pi} |\cos\theta|^q d\theta$ for $1 \leq q \leq p$ (and note that $\frac{1}{2} = \frac{1}{2\pi}\int_0^{2\pi}\cos^2\theta d\theta \leq P_q \leq \frac{1}{2\pi}\int_0^{2\pi}|\cos\theta|d\theta = \frac{2}{\pi}$). Then the embedding $F_{\phi,s} : S \to L_q^1$ (i.e., into a single coordinate of L_q) for vector $v \in S$ is defined by

$$F_s(v) = F_{\phi,s}(v) = \frac{s}{2P_q^{1/q}} \sin\left(\phi + \frac{2}{s}\sum_{i=1}^{m} g_i v_i\right). \tag{1}$$

Note that $0 \leq F_s(v) < s$. For vectors $v, w \in S$ we have that

$$\begin{aligned} |F_s(v) - F_s(w)|^q &= \frac{s^q}{2^q P_q}\left|\sin\left(\phi + \frac{2}{s}\sum_{i=1}^{m} g_i v_i\right) - \sin\left(\phi + \frac{2}{s}\sum_{i=1}^{m} g_i w_i\right)\right|^q \\ &= \frac{s^q}{P_q}\left|\sin\left(\frac{1}{s}\sum_{i=1}^{m} g_i(v_i - w_i)\right)\cos\left(\phi + \frac{1}{s}\sum_{i=1}^{m} g_i(v_i + w_i)\right)\right|^q. \end{aligned}$$

Now, when ϕ is a random variable chosen uniformly from the range $[0, 2\pi]$ and independently of set G, we have that

[3] It may be possible to replace our embedding with that of Mendel and Naor [MN04], but one would still require a tighter analysis, and the final dimension would likely increase. Note also that our embedding is into the reals, while [MN04] embed into the complex numbers; as ℓ_p^d over $\mathbb{C}$ embeds into $\ell_p^{O(\epsilon^{-2}\sqrt{p}2^{p/2}d)}$ over $\mathbb{R}$ with distortion $1+\epsilon$ (a consequence of Dvoretzky's theorem [MS86]), the embedding of [MN04] can in fact be used to achieve an embedding into the reals with increased dimension.

$$\begin{aligned}
\mathbb{E}[|F_s(v) - F_s(w)|^q] &= \frac{s^q}{P_q}\mathbb{E}\left[\left|\sin\left(\frac{1}{s}\sum_{i=1}^m g_i(v_i - w_i)\right)\cos\left(\phi + \frac{1}{s}\sum_{i=1}^m g_i(v_i + w_i)\right)\right|^q\right] \\
&= \frac{s^q}{P_q}\mathbb{E}\left[\left|\sin\left(\frac{1}{s}\sum_{i=1}^m g_i(v_i - w_i)\right)\cos(\phi)\right|^q\right] \\
&= \frac{s^q}{P_q}\mathbb{E}\left[\left|\sin\left(\frac{1}{s}\sum_{i=1}^m g_i(v_i - w_i)\right)\right|^q\right] \cdot \mathbb{E}\left[|\cos(\phi)|^q\right] \\
&= s^q\mathbb{E}\left[\left|\sin\left(\frac{1}{s}\sum_{i=1}^m g_i(v_i - w_i)\right)\right|^q\right]
\end{aligned}$$

where the second equality follows from the periodicity of the cosine function, and the independence of ϕ and G.

This is our single-coordinate embedding. In Lemma 3.1 below, we will describe its behavior on interpoint distances – that is, we derive useful bounds on $\mathbb{E}[|F_s(v) - F_s(w)|^q$. Now $\frac{1}{s}\sum_{i=1}^m g_i(v_i - w_i)$ is distributed as $g\frac{\|v-w\|_p}{s}$ for random variable $g \in G$, so we will set $a = \frac{\|v-w\|_p}{s}$ and will derive bounds for $H(a) = \mathbb{E}[|\sin(ag)|^q] = s^{-q}\mathbb{E}[|F_s(v) - F_s(w)|^q]$. But first we introduce the full embedding f, which is a scaled concatenation of k single-coordinate embeddings: Independently for each coordinate i, create a function $F_{\phi_i} = F_{\phi_i,s}$ by fixing a random angle $0 \le \phi_i \le 2\pi$ and a family of k i.i.d. symmetric p-stables. Then coordinate $f(v)_i$ is defined by

$$f(v)_i = k^{-1/p}F_{\phi_i,s}(v)$$

which can be constructed in $O(m)$ time per coordinate. Note that for $t = \|v - w\|_p$, $\mathbb{E}[|f(v)_i - f(w)_i|^q] = \frac{1}{k}\mathbb{E}[|F_{\phi_i}(v) - F_{\phi_i}(w)|^q] = \frac{s^q}{k}H(t/s)$, and so

$$\mathbb{E}[\|f(v) - f(w)\|_p^q] = s^q H(t/s).$$

Below, we will show that with high probability f is is transform-preserving with respect to H with low distortion.

3.2 Analysis of basic embedding

We now show that both the single-coordinate embedding and the full embedding have desirable properties. Recall that $h(u)$ is the density function of p-stable random variables. Set

$$Q = 2\int_0^\infty u^q h(u)du \approx \frac{1}{p-q}.$$

Also, for $a < 1$ and some fixed $a^2 < \varepsilon < 1$, set

$$Q_a = \frac{1}{2}\int_0^{\sqrt{\varepsilon}/a} u^p h(u)du = \frac{1}{2}\left[\int_0^1 u^p h(u)du + \Theta(\ln(\sqrt{\varepsilon}/a))\right] \approx 1 + \ln(\sqrt{\varepsilon}/a).$$

Lemma 3.1. *Let g be a symmetric p-stable random variable. For $1 \le q \le p \le 2$ and any fixed $0 < \varepsilon < \frac{1}{2}$, $H(a) = \mathbb{E}[|\sin(ag)|^q]$ obeys the following:*

(a). Threshold: $H(a) \leq 1$.

(b). Bi-Lipschitz for small scales: When $q < p$ *and* $a \leq \min\{\varepsilon^{\frac{1}{2}+\frac{1}{p-q}}, \sqrt{\varepsilon}(1+(p-q)\varepsilon^{-(q/2+1)})^{-1/(p-q)}\}$, *we have* $1 - O(\varepsilon) \leq \frac{H(a)}{a^q Q} \leq 1 + O(\varepsilon)$. *When* $q = p$ *and* $a \leq \sqrt{\varepsilon}e^{-\varepsilon^{-(\frac{q}{2}+1)}}$, *we have* $1 - O(\varepsilon) \leq \frac{H(a)}{a^q Q_a} \leq 1 + O(\varepsilon)$.

(c). Bi-Lipschitz for intermediate scales: When $q < p$ *and* $a < 1$, *let* $\delta = 1 - a^{p-q}$), *and we have* $H(a) = \Theta(1 + \delta Q)a^q$. *When* $q = p$ *we have* $H(a) = \Theta\left(1 + \ln(1/a)\right) a^q$.

(d). Lower bound for large scales: When $a \geq 1$, $H(a) > \frac{1}{8}$.

(e). Smoothness: When $a \leq 1$, $\frac{|H((1+\varepsilon)a) - H(a)|}{H(a)} = O(\varepsilon)$.

It follows that the distance transform implied by our single-coordinate embedding F_s achieves the bounds of Lemma 3.1 scaled by s^q, if only in expectation. Item (b) implies that for very small a (i.e., when the inter-point distance under consideration is sufficiently small with respect to the parameter s) the embedding has very small expected distortion (at least when $q < p$). Weaker distortion bounds hold for distances smaller than s (item (c)). For distances greater than s, these are contracted to the threshold (item (a)) or slightly smaller (item (d)). The smoothness property will be useful for constructing the snowflake in Section 4.3. We proceed to consider the full embedding:

Theorem 3.2. *Let* $1 \leq q \leq p \leq 2$ *and* $0 < \varepsilon < \frac{1}{2}$, *and consider an* n*-point set* $S \in l_p^m$. *Set threshold* $s > 1$. *Then with constant probability the oblivious embedding* $f : S \to l_q^k$ *for* $k = O\left(\frac{\log n}{\varepsilon^2} \cdot \min\left\{s^{2q}, \max\left\{\frac{s^{2q-p}}{2q-p}, \varepsilon s^q\right\}\right\}\right)$, *satisfies the following for each point pair* $v, w \in S$, *where* $t = \|v - w\|_p$:

(a). Threshold: $\|f(v) - f(w)\|_q^q < s^q$.

(b). Bi-Lipschitz for large scales: When $t \geq 1$, *we have*

$$(1 - \varepsilon)s^q H(t/s) \leq \|f(v) - f(w)\|_q^q \leq (1 + \varepsilon)s^q H(t/s).$$

(c). Bounded expansion for small scales: When $t < 1$, *we have*

$$\|f(v) - f(w)\|_q^q \leq s^q H(1/s) + \varepsilon.$$

The embedding can be constructed in $O(mk)$ *time per point.*

Theorem 3.2 demonstrates that in a certain range, there exists a transform-preserving embedding with high probability. (We note that when $2q$ is close to p, better dimension bounds can be obtained by embedding into an interim value $q + c$ for some c and then embedding into q.)

Proof of Lemma 3.1. **Item (a).** This follows trivially from the fact that $|\sin(x)| \leq 1$.
Item (b). Note that since the density function h is symmetric about 0,

$$H(a) = 2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du + 2\int_{\sqrt{\varepsilon}/a}^{\infty} |\sin(au)|^q h(u)du. \tag{2}$$

We show that under the conditions of the item, the second term in Equation (2) is dominated by the first. Considering the second term, we have that

$$2\int_{\sqrt{\varepsilon}/a}^{\infty} |\sin(au)|^q h(u)du < 2\int_{\sqrt{\varepsilon}/a}^{\infty} h(u)du < 2c'_p \int_{\sqrt{\varepsilon}/a}^{\infty} \frac{1}{u^{p+1}}du = \frac{2c'_p}{p}\left(\frac{a}{\sqrt{\varepsilon}}\right)^p. \quad (3)$$

Turning to the first term, recall the Taylor series expansion $\sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \ldots$; so when $x < \sqrt{\varepsilon}$ we have that $x(1 - \frac{\varepsilon}{6}) < \sin(x) < x$. Also note that the conditions of the item give that $a < \sqrt{\varepsilon}$, and so $\frac{\sqrt{\varepsilon}}{a} > 1$. Hence, when $q < p$ we have

$$\begin{aligned}
2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du &> 2(1-\varepsilon/6)^q a^q \int_0^{\sqrt{\varepsilon}/a} u^q h(u)du \\
&> 2(1-\varepsilon/3)a^q c_p \left[\int_0^1 \frac{u^q}{1+u^{p+1}}du + \int_1^{\sqrt{\varepsilon}/a} \frac{u^q}{1+u^{p+1}}du\right] \\
&> 2(1-\varepsilon/3)a^q c_p \left[\int_0^1 \frac{u^q}{2}du + \int_1^{\sqrt{\varepsilon}/a} \frac{u^{q-p-1}}{2}du\right] \\
&= (1-\varepsilon/3)a^q c_p \left[\frac{1}{q+1} + \frac{1-(a/\sqrt{\varepsilon})^{p-q}}{p-q}\right]
\end{aligned}$$

When $q = p$, the same analysis gives

$$2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du > (1-\varepsilon/3)a^q c_p \left[\frac{1}{q+1} + \ln(\sqrt{\varepsilon}/a)\right]$$

We proceed to show that the first term of Equation (2) exceeds the second by a factor of $\Omega(\varepsilon^{-1})$. For $q = p$ this holds trivially when $a^q \ln(\sqrt{\varepsilon}/a) \geq \frac{1}{\varepsilon}\left(\frac{a}{\sqrt{\varepsilon}}\right)^p$ – or equivalently, when $\ln(\sqrt{\varepsilon}/a) \geq \varepsilon^{-(\frac{q}{2}+1)}$ – which in turn holds exactly when $a \leq \sqrt{\varepsilon}e^{-\varepsilon^{-(\frac{q}{2}+1)}}$. For $q < p$, the condition is fulfilled when $a^q \geq \frac{1}{\varepsilon}\left(\frac{a}{\sqrt{\varepsilon}}\right)^p$, which holds exactly when $a \leq \varepsilon^{\frac{p/2+1}{p-q}}$. Better, the condition is also fulfilled when $a^q\left[\frac{1-(a/\sqrt{\varepsilon})^{p-q}}{p-q}\right] \geq \frac{1}{\varepsilon}\left(\frac{a}{\sqrt{\varepsilon}}\right)^p$, and this is equivalent to satisfying $a^{p-q}\left[\frac{p-q}{\frac{p}{2}+1} + \frac{1}{\varepsilon^{(p-q)/2}}\right] \leq 1$; this holds exactly when $a \leq \sqrt{\varepsilon}(1+(p-q)\varepsilon^{-(q/2+1)})^{-1/(p-q)}$.

As the first term of Equation (2) dominates the second, it follows that when $q < p$ then for some constant c

$$\begin{aligned}
H(a) &= 2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du + 2\int_{\sqrt{\varepsilon}/a}^{\infty} |\sin(au)|^q h(u)du \\
&\leq 2(1+c\varepsilon)\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du \\
&< 2(1+c\varepsilon)a^q \int_0^{\infty} u^q h(u)du \\
&= (1+c\varepsilon)a^q Q.
\end{aligned}$$

Further, we have that

$$\begin{aligned} H(a) &= 2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du + 2\int_{\sqrt{\varepsilon}/a}^{\infty} |\sin(au)|^q h(u)du \\ &> 2\int_0^{\sqrt{\varepsilon}/a} |\sin(au)|^q h(u)du \\ &> 2(1-\varepsilon/3)a^q \int_0^{\sqrt{\varepsilon}/a} u^q h(u)du \\ &> (1-\varepsilon/3)(1-c'\varepsilon)a^q Q, \end{aligned}$$

where the final inequality follows from noting that since $a \leq \varepsilon^{\frac{1}{2}+\frac{1}{p-q}}$, $\int_{\sqrt{\varepsilon}/a}^{\infty} u^q h(u)du \leq \int_{\varepsilon^{-1/(p-q)}}^{\infty} u^q h(u)du \approx \frac{\varepsilon}{p-q} \approx \varepsilon Q$, and so $\int_0^{\sqrt{\varepsilon}/a} u^q h(u)du = \int_0^{\infty} u^q h(u)du - \int_{\sqrt{\varepsilon}/a}^{\infty} u^q h(u)du > (1-c'\varepsilon)Q$ for some c'.

This completes the analysis for $q < p$. The same analysis gives that when $q = p$,

$$H(a) < 2(1+c\varepsilon)a^q \int_0^{\sqrt{\varepsilon}/a} u^q h(u)du = (1+c\varepsilon)a^q Q_a,$$

and

$$H(a) > 2(1-\varepsilon/3)a^q \int_0^{\sqrt{\varepsilon}/a} u^q h(u)du = (1-\varepsilon/3)a^q Q_a.$$

Item (c). The analysis is similar to that presented in the proof of Item (b). Noting that when $0 \leq x \leq 1$, $\sin(x) \approx x$, and recalling that under the conditions of the item $a \leq 1$, we have for $q < p$ that

$$\begin{aligned} H(a) &= 2\int_0^{1/a} |\sin(au)|^q h(u)du + 2\int_{1/a}^{\infty} |\sin(au)|^q h(u)du. \\ &= \Theta\left(a^q \int_0^{1/a} u^q h(u)du\right) + O\left(\int_{1/a}^{\infty} h(u)du\right) \\ &= \Theta\left(\left(1+\frac{1-a^{p-q}}{p-q}\right)a^q\right) + O(a^p) \\ &= \Theta\left(\left(1+\frac{\delta}{p-q}\right)a^q\right) + O(a^p) \\ &= \Theta\left(\left(1+\frac{\delta}{p-q}\right)a^q\right). \end{aligned}$$

Similarly, for $q = p$ we have

$$H(a) = \Theta\left(a^q \int_0^{1/a} u^p h(u)du\right) + O\left(\int_{1/a}^{\infty} h(u)du\right) = \Theta\left(a^q \int_0^{1/a} u^p h(u)du\right) \approx (1+\ln(1/a))\, a^p.$$

Item (d). First note that when $p \geq 1$, $h(x) = \frac{1}{\pi}\int_0^{\infty} \cos(tx)e^{-t^p}dt < \frac{1}{\pi}\int_0^{\infty} e^{-t^p}dt < \frac{1}{\pi}$, so $\int_0^b h(u)du < \frac{b}{\pi}$. Since $h(x)$ is a symmetric density function, we have $\int_0^{\infty} h(u)du = \frac{1}{2}$, and so

$\int_b^\infty h(u)du > \frac{1}{2} - \frac{b}{\pi}$. We have for any $0 < \theta < \frac{\pi}{2}$,

$$\begin{aligned}
H(a) &\geq 2\int_{\theta/a}^{\infty} |\sin(au)|^q h(u)du \\
&> 2\sum_{i=0}^{\infty} \int_{\frac{i\pi+\theta}{a}}^{\frac{(i+1)\pi-\theta}{a}} |\sin(au)|^q h(u)du \\
&> 2|\sin(\theta)|^q \sum_{i=0}^{\infty} \int_{\frac{i\pi+\theta}{a}}^{\frac{(i+1)\pi-\theta}{a}} h(u)du \\
&> 2|\sin(\theta)|^q \left[1 - \frac{2\theta}{\pi}\right] \sum_{i=0}^{\infty} \int_{\frac{i\pi+\theta}{a}}^{\frac{(i+1)\pi+\theta}{a}} h(u)du \\
&= 2|\sin(\theta)|^q \left[1 - \frac{2\theta}{\pi}\right] \int_{\theta/a}^{\infty} h(u)du \\
&> |\sin(\theta)|^q \left[1 - \frac{2\theta}{\pi}\right] \left[1 - \frac{2\theta}{a\pi}\right].
\end{aligned}$$

Where the fourth inequality follows from the fact that $h(x)$ is monotone decreasing for $x \geq 0$. The claimed result follows by taking $a = 1$ (the maximum value of a) and $\theta = \frac{\pi}{4}$, and recalling that $q \leq 2$.

Item (e). As in the proof of Item (c) above, we have that for $q \leq p$ and $a \leq 1$, $H(a) \approx \int_0^{1/a} |\sin(au)|^q h(u)du + \int_{1/a}^\infty h(u)du = \int_0^{1/a} |\sin(au)|^q h(u)du + O(a^p) \approx \int_0^{1/a} |\sin(au)|^q h(u)du$. Now recall the Taylor series expansion $\cos(x) = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \ldots > 1 - \frac{x^2}{2}$ (when $0 \leq x \leq 1$), and note that as a consequence of the Mean Value Theorem, $||A|^q - |B|^q| = qC^{q-1}||A| - |B||$ for some $|A| \geq C \geq |B|$. Further noting that when $u \leq \frac{1}{a}$ we have $au \leq 1$ and that when $u \leq \frac{1}{\varepsilon a}$ we have $\varepsilon au \leq 1$, we conclude that

$$\begin{aligned}
H(a(1+\varepsilon)) - H(a) &= 2\int_0^\infty [|\sin(a(1+\varepsilon)u)|^q - |\sin(au)|^q] h(u)du \\
&\leq 2\int_0^\infty q[\max\{|\sin(a(1+\varepsilon)u)|, |\sin(au)|\}]^{q-1} \cdot \\
&\quad [||\sin(a(1+\varepsilon)u)| - |\sin(au)||] h(u)du \\
&\leq 2\int_0^\infty q[\max\{|\sin(au)\cos(\varepsilon au)| + |\sin(\varepsilon au)\cos(au)|, |\sin(au)|\}]^{q-1} \cdot \\
&\quad [||\sin(au)\cos(\varepsilon au)| + |\sin(\varepsilon au)\cos(au)| - |\sin(au)||] h(u)du \\
&\leq 2\int_0^\infty q[|\sin(au)|^{q-1} + |\sin(\varepsilon au)|^{q-1}][|\sin(\varepsilon au)|] h(u)du \\
&= O\left(\int_0^{1/a} \varepsilon |\sin(au)|^q h(u)du + \varepsilon a \int_{1/a}^{1/(\varepsilon a)} uh(u)du + \int_{1/(\varepsilon a)}^\infty h(u)du\right) \\
&= O\left(\int_0^{1/a} \varepsilon |\sin(au)|^q h(u)du + \varepsilon a^p + \varepsilon^p a^p\right) \\
&= O(\varepsilon H(a))
\end{aligned}$$

□

Proof of Theorem 3.2. The first claim of the theorem follows from the fact that for all i, $0 \leq F_{\phi_i,s}(v) < s$.

To prove the rest of the theorem, we may make use of Hoeffding's inequality. When $1 \leq t \leq s$, we have by Lemma 3.1(b)(c)(d) that $s^q H(t/s) = s^q \Omega(1) = \Omega(1)$. Then Claim A.1 implies that for some $k = O\left(\frac{s^{2q}}{\varepsilon^2} \log n\right)$, we have

$$\begin{aligned} \Pr\left[|\|f(v) - f(w)\|_q^q - s^q H(t/s)| > \varepsilon s^q H(t/s)\right] &= \Pr\left[|\|f(v) - f(w)\|_q^q - s^q H(t/s)| > \varepsilon \Omega(1)\right] \\ &\leq \frac{1}{n^2}, \end{aligned}$$

so this distortion guarantee can hold simultaneously for all point pairs. When $t < 1$, first note that when $H(t/s) = \Theta(1)$, we can use the same proof as for $1 \leq t \leq s$ above. If $H(t/s) = o(1)$, we have

$$\begin{aligned} \Pr\left[\|f(v) - f(w)\|_q^q > s^q H(1/s) + \varepsilon\right] &= \Pr\left[\|f(v) - f(w)\|_q^q - s^q H(t/s) > s^q H(1/s) - s^q H(t/s) + \varepsilon\right] \\ &< \Pr\left[\|f(v) - f(w)\|_q^q - s^q H(t/s) > \varepsilon\right] \\ &\leq \frac{1}{n^2}. \end{aligned}$$

An alternate bound can be derived using Bennett's inequality (Claim A.2). For this, it suffices to take $k = O\left(\frac{s^{2q} \log n}{\sigma^2 V\left(\frac{s^q \varepsilon [s^q H(t/s)]}{\sigma^2}\right)}\right)$, with the variance term $\sigma^2 = \Theta\left(s^{2q} \mathbb{E}[|\sin(ag)|^{2q}]\right)$. We will prove the case $t \geq 1$, and the case $t < 1$ follows as above. Set $r = \frac{s^q \varepsilon [s^q H(t/s)]}{\sigma^2}$. If $r \geq 1$, then we have $V(r) = \Omega(r)$, and so $k = O\left(\frac{s^q \log n}{\varepsilon [s^q H(t/s)]}\right) = O\left(\frac{s^q \log n}{\varepsilon}\right)$. Otherwise $r < 1$, and we have $V(r) = \Theta(r^2)$, from which we derive $k = O\left(\frac{\sigma^2 \log n}{\varepsilon^2 [s^q H(t/s)]^2}\right)$. Now, if $t \geq s$, we recall by 3.1(d) that $H(t/s) = \Theta(1)$, and noting that $\sigma^2 \approx s^{2q} \mathbb{E}[|\sin(ag)|^{2q}] = O(s^{2q})$ we obtain $k = O\left(\frac{\log n}{\varepsilon^2}\right)$. When $1 \leq t < s$, we have by 3.1(b)(c) that $H(t/s) = \Omega((t/s)^q)$, and so $k = O\left(\frac{\sigma^2 \log n}{\varepsilon^2 t^{2q}}\right)$. In this case we require a better bound on σ^2: Setting $a = \frac{t}{s} < 1$ we have (by analysis similar to the proof of Lemma 3.1)

$$\begin{aligned} \sigma^2 &\approx s^{2q} \mathbb{E}[|\sin(ag)|^{2q}] \\ &\approx s^{2q} \left[a^{2q} \int_0^1 u^{2q} du + a^{2q} \int_1^{1/a} u^{2q-p-1} du + \int_{1/a}^{\infty} u^{-p-1} du\right] \\ &\approx s^{2q} \left[\frac{a^{2q}}{2q+1} + \frac{a^p}{2q-p} + \frac{a^p}{p}\right] \\ &\approx s^{2q} \frac{a^p}{2q-p}. \end{aligned}$$

It follows that $k = O\left(\frac{a^{p-2q} \log n}{(2q-p)\varepsilon^2}\right) = O\left(\frac{s^{2q-p} \log n}{(2q-p)\varepsilon^2}\right)$. □

4 Range and snowflake embeddings

In Section 3 above, we presented our basic embedding. Here, we show how to use the basic embedding to derive a dimension-reducing embedding that preserves distances in a fixed range (Section 4.1). We also show that the basic embedding can be used to derive a dimension-reducing

snowflake embedding (Section 4.3). As a precursor to the snowflake embedding, we show that the basic and range embeddings can be improved to embed into the doubling dimension of the space (Section 4.2).

4.1 Range embedding

By combining Theorem 3.2 and Lemma 3.1 and choosing an appropriate parameter s, we can achieve a dimension-reducing oblivious strong range embedding for ℓ_p. This is the central contribution of our paper:

Theorem 4.1. *Let $1 \le q \le p \le 2$ and $0 < \varepsilon < \frac{1}{2}$, and consider an n-point set $S \in l_p^m$. Fix range $R > 1$ and set threshold s as follows: When $q < p$, $s \approx R\varepsilon^{-1/2}(1+(p-q)\varepsilon^{-(q/2+1)})^{1/(p-q)}$, and when $q = p$, $s \approx \max\{R^{1/\varepsilon}, R\varepsilon^{-1/2}e^{\varepsilon^{-(\frac{q}{2}+1)}}\}$. Then there exists an oblivious embedding $f : S \to l_q^k$ for $k = O\left(\frac{\log n}{\varepsilon^2} \cdot \min\left\{s^{2q}, \max\left\{\frac{s^{2q-p}}{2q-p}, \varepsilon s^q\right\}\right\}\right)$, which satisfies the following for each point pair $v, w \in S$, where $t = \|v - w\|_p$:*

(a). Threshold: When $q < p$ we have $\|f(v) - f(w)\|_q^q \le \frac{s^q}{Q}$.

When $q = p$ we have $\|f(v) - f(w)\|_q^q \le \frac{s^q}{Q_{R/s}}$.

(b). Bounded expansion and contraction for large scales: When $t > R$, we have $\|f(v) - f(w)\|_q^q = O(t^q)$, and $\|f(v) - f(w)\|_q^q \ge R^q + \varepsilon$.

(c). Bi-Lipschitz for intermediate scales: When $1 \le t \le R$, we have $(1-\varepsilon)t^q \le \|f(v) - f(w)\|_q^q \le (1+\varepsilon)t^q$.

(d). Bounded expansion for small scales: When $t < 1$, we have $\|f(v) - f(w)\|_q^q \le 1 + \varepsilon$.

The embedding can be constructed in $O(mk)$ time per point.

Proof. We use the construction of Theorem 3.2 with the stated value of s, and then scale down by a factor of $Q^{1/q}$ or $Q_{R/s}^{1/q}$. (Note that $Q, Q_{R/s} = \Omega(1)$.) Then the threshold guarantee follows immediately from Theorem 3.2(a) and the scaling step.

We will now prove the rest of the theorem for the case $p < q$. First, the bi-Lipschitz claim for values $1 \le t \le R$ follows immediately from Theorem 3.2(b) and Lemma 3.1(b), when noting that for an appropriate choice of s, $\frac{t}{s} \le \varepsilon^{1/2}(1+(p-q)\varepsilon^{-(q/2+1)})^{-1/(p-q)}$. Similarly, the bounded expansion claim for $t < 1$ follows from Theorem 3.2(c) and the aforementioned bi-Lipschitz guarantee at $t = 1$. Finally, the bounded expansion for $t > R$ follows from Theorem 3.2(b), and the bounded contraction follows from Theorem 3.2(b) combined with fact that when $a > R/s$, $H(a) > H(R/s)$ (as a consequence of Lemma 3.1(c)(d) for an appropriate choice of s).

For $p = q$, we have essentially the same proof, only noting that the function $Q_a = Q_{t/s}$ is monotone decreasing in t, and since $s = O(R^{1/\varepsilon})$ we have for all $1 \le t \le R$ that $\frac{Q_{t/s}}{Q_{R/s}} \le \frac{Q_{1/s}}{Q_{R/s}} = O\left(\frac{1+\varepsilon^{-1}\log R}{1+(\varepsilon^{-1}-1)\log R}\right) = O(1+\varepsilon)$. ∎

4.2 Intrinsic dimensionality reduction

Here we show that the guarantees of Theorem 3.2 and Theorem 4.1 can be achieved by (non-oblivious) embeddings whose target dimension in independent of n, and depends only on the doubling dimension of the space. The following lemma is derived by applying the framework of [GK11] to Theorem 3.2.

Lemma 4.2. *Let $1 \le q \le p \le 2$ and $0 < \varepsilon < \frac{1}{2}$, and consider an n-point set $S \in l_p^m$. Set threshold $s > 1$. Then there exists an embedding $f : S \to l_q^k$ for $k = \tilde{O}\left(\frac{\operatorname{ddim}^2(S)}{\varepsilon^3} \cdot s \cdot \min\left\{s'^{2q}, \max\left\{\frac{s'^{2q-p}}{2q-p}, \varepsilon s'^q\right\}\right\}\right)$ for $s' = \frac{s \operatorname{ddim}(S)}{\varepsilon}$, which satisfies the following for each point pair $v, w \in S$, where $t = \|v - w\|_p$:*

(a). Threshold: $\|f(v) - f(w)\|_q^q < s^q$.

(b). Bi-Lipschitz for intermediate scales: When $1 \le t \le s$, we have

$$(1-\varepsilon)s^q H(t/s) \le \|f(v) - f(w)\|_q^q \le (1+\varepsilon)s^q H(t/s).$$

(c). Strong bounded expansion for small scales: When $t < 1$, for some constant c

$$\|f(v) - f(w)\|_q^q \le \min\{(1+\varepsilon)s^q H(1/s), c \operatorname{ddim}(S) t\}.$$

Given a point hierarchy for S, the embedding can be constructed in $2^{\tilde{O}\operatorname{ddim}(S)} + O(mk)$ time per point.

Proof. Similar to what was done in [GK11], we compute for S a padded decomposition with padding s. This is a multiset $[P_1, \ldots, P_m]$ where each partition P_i is a set of clusters, and every point is s-padded in a $\left(1 - \frac{\varepsilon}{s}\right)$ fraction of the partitions. Each cluster has diameter bounded by $O\left(s \operatorname{ddim}(S)\right)$, and the support is $m = \tilde{O}(s\varepsilon^{-1} \operatorname{ddim}(S))$. Using the hierarchy, this can be done in time $2^{\tilde{O}(\operatorname{ddim}(S))}$ per point [BGK$^+$11].

We embed each partition P_i separately as follows: For each cluster $C \in P_i$, we extract from C an $\frac{\varepsilon}{\operatorname{ddim}(S)}$-net $N \subset C$ Now each cluster net has aspect ratio $\frac{s \operatorname{ddim}^2(S)}{\varepsilon}$, and so $|N| = \left(\frac{s}{\varepsilon}\right)^{\tilde{O}\operatorname{ddim}(S)}$. We then scale N up by a factor of $\frac{\operatorname{ddim}(S)}{\varepsilon}$ so that the minimum distance is 1, invoke the embedding of Theorem 3.2 with parameter s' on N, and scale back down. This procedure has a runtime cost of $O(mk)$ per point, thresholds all distances at s, and reduces dimension to $\tilde{O}\left(\frac{\operatorname{ddim}(S)}{\varepsilon^2} \cdot \min\left\{s'^{2q}, \max\left\{\frac{s'^{2q-p}}{2q-p}, \varepsilon s'^q\right\}\right\}\right)$. We then concatenate the m cluster partitions together and scale down by $m^{1/q}$, achieveing an embedding of dimension k for the net points. We then extend this embedding to all points using the $c \operatorname{ddim}(S)$-factor Lipschitz extension of Lee and Naor [LN05] for metric space.

For the net points, item (a) holds immediately. For item (b), we note that when $1 \le t \le s$, the points fall in the same cluster in a fraction $\left(1 - \frac{\varepsilon}{s}\right)$ of the partitions, and these partitions contribute the correct amount to the interpoint distance. However, in a fraction $\frac{\varepsilon}{s}$ of the partition the points are not found in the same cluster, and in these cases the contribution may be as large as s. These partitions account for an additive value of at most $\frac{\varepsilon}{s} \cdot s = \varepsilon \le \varepsilon t$. Item (c) follows in an identical manner.

For the non-net points, item (a) holds since we may assume that all interpoint distance are at most s, since the non-net points outside the convex hull of the net-points can all be projected

onto the hull, and this can only improve the quality of the extension. Item (b)(c) follow from the embedding of the $\frac{\varepsilon}{\mathrm{ddim}(S)}$-net: By the guarantees of the extension, an embedded non-net point may be at distance at most $O(\varepsilon)$ from its closest net-point, and then the items follow by an appropriate scaling down of ε. □

Similarly, the exact guarantees of Theorem 4.1 can be achieved by a non-oblivious embedding with dimension independent of n.

Corollary 4.3. *Let $1 \le q \le p \le 2$ and consider an n-point set $S \in l_p^m$. Set threshold $s > 1$. Fix range $R > 1$ and set threshold s as follows: When $q < p$, $s \approx R\varepsilon^{-1/2}(1+(p-q)\varepsilon^{-(q/2+1)})^{1/(p-q)}$, and when $q = p$, $s \approx \max\{R^{1/\varepsilon}, R\varepsilon^{-1/2}e^{\varepsilon^{-(\frac{q}{2}+1)}}\}$. Then there exists an embedding $f : S \to l_q^k$ satisfying items (a)(c)(d) of Theorem 4.1. The target dimension is $k = \tilde{O}\left(\frac{\mathrm{ddim}^2(S)}{\varepsilon^3} \cdot s \cdot \min\left\{s'^{2q}, \max\left\{\frac{s'^{2q-p}}{2q-p}, \varepsilon s'^q\right\}\right\}\right)$ for $s' = \frac{s\,\mathrm{ddim}(S)}{\varepsilon}$. Given a point hierarchy for S, the embedding can be constructed in $2^{\tilde{O}\,\mathrm{ddim}(S)} + O(mk)$ time per point.*

Comment. We conjecture that the $\frac{\mathrm{ddim}^2(S)}{\varepsilon^3}$ terms can be reduced to $\frac{\mathrm{ddim}(S)}{\varepsilon^2}$ by combining the randomness used separately for the construction of the padded decomposition, threshold embedding f, and the Lipschitz extension (as was done in [BRS11] for ℓ_2).

4.3 Snowflake embedding

The embedding of Lemma 4.2 implies a global dimension-reduction snowflake embedding for ℓ_p.

Lemma 4.4. *Let $0 < \varepsilon < 1/4$, $0 < \alpha < 1$, and $\tilde{\alpha} = \min\{\alpha, 1-\alpha\}$. Every finite subset $S \subset \ell_p$ admits an embedding $\Phi : S \to \ell_q^k$ ($1 \le q \le p \le 2$) for $k = \tilde{O}\left(\frac{\mathrm{ddim}^6(S)}{\tilde{\alpha}^2\varepsilon^8} \cdot \min\left\{s'^{2q}, \max\left\{\frac{s'^{2q-p}}{2q-p}, \varepsilon s'^q\right\}\right\}\right)$ where $s' = \left(\frac{\mathrm{ddim}(S)}{\varepsilon}\right)^4$ and $1 \le \frac{\|\Phi(x)-\Phi(y)\|_q}{\|x-y\|_p^\alpha} \le 1+\varepsilon$, $\quad \forall x, y \in S$. Given a point hierarchy for S, the embedding can be constructed in $2^{\tilde{O}\,\mathrm{ddim}(S)} + O(mk)$ time per point.*

To prove Lemma 4.4, we will make use of the snowflake framework presented in [GK11]. For simplicity, we will first prove the theorem for $\alpha = 1/2$, and then extend the proof to arbitrary $0 < \alpha < 1$. Fix a finite set $S \subset \ell_p$, and assume without loss of generality that the minimum interpoint distance in S is 1. Define $v = 2\left\lceil \log_{1+\varepsilon}\left(\frac{\mathrm{ddim}(S)}{\varepsilon}\right)\right\rceil = O\left(\frac{1}{\varepsilon}\log\frac{\mathrm{ddim}(S)}{\varepsilon}\right)$, and the set $I = \{i \in \mathbb{Z} : \ (1+\varepsilon)^{-2v} \le (1+\varepsilon)^i \le (1+\varepsilon)^{2v}\,\mathrm{diam}(S)\}$. For each $i \in I$, fix $r_i = (1+\varepsilon)^i$, and let $\varphi_i : S \to \ell_p^{k'}$ for $k' = O(k\varepsilon)$ be the embedding that achieves the bounds of Lemma 4.2 with parameter $s = (1+\varepsilon)^{2v} = \left(\frac{\mathrm{ddim}(S)}{\varepsilon}\right)^4$ and scaled by $\frac{r_i}{\sqrt{s}}$ so that the bi-Lipschitz property Lemma 4.2(b) holds whenever $\frac{r_i}{\sqrt{s}} \le t \le \sqrt{s}r_i$ and the threshold guarantee of Lemma 4.2(a) holds at $\sqrt{s}r$.

We shall now use the following technique due to Assouad [Ass83]. First, each φ_i is scaled down by $1/\sqrt{r_i} = (1+\varepsilon)^{-i/2}$. They are then grouped in a round-robin fashion into $2v$ groups, and the embeddings in each group are summed up. This yields $2v$ embeddings, each into $\ell_p^{k'}$; these are combined using a direct-sum, resulting in one map Φ into $\ell_p^{2vk'}$. Formally, let $i \equiv_v j$ denote that two integers i, j are equal modulo v. Define $\Phi : S \to \ell_p^{2vk'}$ using the direct sum $\Phi = \bigoplus_{j\in[2v]} \Phi_j$,

where each $\Phi_j : S \to \ell_p^{k'}$ is given by

$$\Phi_j = \sum_{i \in I:\ i \equiv_{2v} j} \frac{\varphi_i}{(1+\varepsilon)^{i/2}}.$$

For $M = M(\varepsilon) > 0$ that will be defined later, our final embedding is $\Phi/\sqrt{M} : S \to \ell_p^{2vk'}$, which has target dimension $2vk' = O(k)$, as required (for $\alpha = 1/2$). It thus remains to prove the distortion bound. Define $B_i = \frac{\|\varphi_i(x) - \varphi_i(y)\|_p}{(1+\varepsilon)^{i/2}}$. We will need the following lemma.

Lemma 4.5. *Let $\Phi : S \to \ell_p^{2vk'}$ be as above, let $x, y \in S$. Then for every interval $A \subset I$ of size $2v$ (namely $A = \{a - v, \ldots, a, \ldots, a + v - 1\}$),*

$$\begin{aligned}
\|\Phi(x) - \Phi(y)\|_p^p &\le \sum_{i \in A} \Big(B_i \ + \sum_{i' \in I \setminus A:\ i' \equiv_{2v} i} B_{i'} \Big)^p \\
&\le \sum_{i \in A} \Big(B_i^p + \ 2B_i \sum_{i' \in I \setminus A:\ i' \equiv_{2v} i} B_{i'} + \Big(\sum_{i' \in I \setminus A:\ i' \equiv_{2v} i} B_{i'} \Big)^p \Big) \\
&\le \sum_{a - v \le i < a} \Big(B_i^p + B_{i+v}^p + \ 2(B_i + B_{i+v}) \sum_{i' \in I \setminus A:\ i' \equiv_v i} B_{i'} + \Big(\sum_{i' \in I \setminus A:\ i' \equiv_v i} B_{i'} \Big)^p \Big) \\
\|\Phi(x) - \Phi(y)\|_p^p &\ge \sum_{i \in A} \Big(\max\Big\{ 0, B_i \ - \sum_{i' \in I \setminus A:\ i' \equiv_{2v} i} B_{i'} \Big\} \Big)^p \\
&\ge \sum_{a - v \le i < a} \Big(B_i^p + B_{i+v}^p - \ 2(B_i + B_{i+v}) \sum_{i' \in I \setminus A:\ i' \equiv_v i} B_{i'} \Big)
\end{aligned}$$

Proof. By construction,

$$\Big\| \Phi(x) - \Phi(y) \Big\|_p^p = \sum_{j \in [v]} \Big\| \Phi_j(x) - \Phi_j(y) \Big\|_p^p = \sum_{i \in A} \Big\| \sum_{i' \in I:\ i' \equiv_{2v} i} \frac{\varphi_{i'}(x) - \varphi_{i'}(y)}{(1+\varepsilon)^{i/2}} \Big\|_p^p.$$

Fix $i \in A$ and let us bound the term corresponding to i. The first required inequality now follows by separating (among all $i' \in I$ with $i' \equiv_{2v} i$) the term for $i' = i$ from the rest, and applying the triangle inequality for vectors $v_1, \ldots, v_s \in \ell_p^{k'}$, namely, $\| \sum_l v_l \|_p \le \sum_l \|v_l\|_p$. We then bound the sum of the terms for indices i and $i + v$. The second inequality follows similarly by separating the term for $i' = i$ from the rest, and applying the following triangle inequality for vectors $u, v_1, \ldots, v_s \in \ell_p^k$, namely, $\|u + \sum_l v_l\|_p \ge \max\{0, \|u\|_p - \sum_l \|v_l\|_p\}$. We then bound the sum of the terms for indices i and $i + v$. (Note that $(\max\{0, b - c\})^p \ge b^p - 2bc$.) ☐

The proof of Lemma 4.4 proceeds by demonstrating that, for an appropriate choice of A (meaning a, having fixed v), the leading terms in the above summations (B_i^p and B_{i+v}^p for $a \le i < a + v$) dominate the sum of all other terms of the summations. Fix $x, y \in S$, let $i^* \in I$ be such that $(1+\varepsilon)^{i^*} \le \|x - y\|_p \le (1+\varepsilon)^{i^*+1}$, and let $a = i^*$ so that $A = \{i^* - v, \ldots, i^* + v - 1\}$. Consider $i \in A$; we have the following lemma:

Lemma 4.6. *The following hold for all $i^* - v \le i < i^*$:*

(a). $\varepsilon \cdot B_i = \Omega\left(\sum_{i' \in I \setminus A:\ i' \equiv_v i, i' < i} B_{i'} \right)$

(b). $\varepsilon \cdot B_{i+v} = \Omega\left(\sum_{i' \in I \setminus A:\ i' \equiv_v i, i' > i} B_{i'}\right)$

(c). $\varepsilon \cdot (B_i + B_{i+v}) = \Omega\left(\sum_{i' \in I \setminus A:\ i' \equiv_v i} B_{i'}\right)$

Proof. We prove each item in turn.

(a) By Lemma 4.2(b) and Lemma 3.1(c), we have that $\|\varphi_i(x) - \varphi_i(y)\|_p = \Omega(r_i)$, from which it follows that $B_i = \Omega(1+\varepsilon)^{j/2}$. Similarly we have that for all $j < i$, $\|\varphi_j(x) - \varphi_j(y)\|_p = O(\log(\mathrm{ddim}(S)/\varepsilon) \cdot r_j)$, from which it follows that $B_j = O((1+\varepsilon)^{j/2})$. Recalling that a geometric series with constant ratio sum to some constant times the largest largest term, we have that

$$\sum_{i' \in I \setminus A:\ i' \equiv_v i, i' < i} B_{i'} = O\left(\log(\mathrm{ddim}(S)/\varepsilon)(1+\varepsilon)^{(i-v)/2}\right) = O\left(\varepsilon(1+\varepsilon)^{i/2}\right) = O\left(\varepsilon B_i\right).$$

(b) By Lemma 4.2(b) and Lemma 3.1(c), we have that $\|\varphi_{i+v}(x) - \varphi_{i+v}(y)\|_p = \Omega(\|x-y\|_p) = \Omega((1+\varepsilon)^{i^*})$, so $B_{i+v} = \Omega\left((1+\varepsilon)^{i^* - (i+v)/2}\right)$. By Lemma 4.2(c), we have for $j > i$ that $\|\varphi_j(x) - \varphi_j(y)\|_p = O\left(\mathrm{ddim}(S)(\|x-y\|_p)\right) = O(\mathrm{ddim}(S)(1+\varepsilon)^{i^*})$. It follows that for $i' \equiv_v i$ and $i' > i$ we have $B_{i'} = O\left(\mathrm{ddim}(S)(1+\varepsilon)^{i^* - i'/2}\right)$, and thus

$$\sum_{i' \in I \setminus A:\ i' \equiv_v i, i' > i} B_{i'} = O\left(\mathrm{ddim}(S)(1+\varepsilon)^{i^* - (i+2v)/2}\right) = O\left(\varepsilon B_{i+v}\right).$$

(c) This follows trivially from parts (a) and (b).

□

Now, plugging Lemma 4.6 into Lemma 4.5, we obtain

$$\begin{aligned}
\|\Phi(x) - \Phi(y)\|_p^p &\geq \sum_{i \in A:\ i < i^*} \left(B_i^p + B_{i+v}^p - p(B_i + B_{i+v}) \sum_{i' \in I \setminus A:\ i' \equiv_v i} B_{i'}\right) \\
&\geq (1 - O(\varepsilon)) \sum_{i \in A:\ i < i^*} \left(B_i^p + B_{i+v}^p\right) \\
&= (1 - O(\varepsilon)) \sum_{i \in A} B_i^p \\
&= (1 - O(\varepsilon)) \sum_{i \in A} \frac{\|\varphi_i(x) - \varphi_i(y)\|_p}{(1+\varepsilon)^{pi/2}} \\
&= (1 - O(\varepsilon)) M \frac{\|x-y\|_p^p}{(1+\varepsilon)^{pi^*/2}},
\end{aligned}$$

for some fixed constant M, where the last step follows from Theorem 3.2 and Lemma 3.1(e)(c). Similarly, $\|\Phi(x) - \Phi(y)\|_p^p \leq (1 - O(\varepsilon)) M \frac{\|x-y\|_p^p}{(1+\varepsilon)^{pi/2}}$. We conclude that the final embedding $\Phi / M^{1/p}$ achieves distortion $1 + O(\varepsilon)$ for $\alpha = 1/2$.

Arbitrary $0 < \alpha < 1$. Turning to proving the theorem for arbitrary values of $0 < \alpha < 1$, we repeat the previous construction and proof with $v = 2^{\frac{\lceil \log_{1+\varepsilon}(\frac{\mathrm{ddim}(S)}{\varepsilon})\rceil}{\tilde{\alpha}}} = O(\frac{1}{\tilde{\alpha}\varepsilon}\log\frac{\mathrm{ddim}(S)}{\varepsilon})$. As before, $\varphi_i : S \to \ell_2^{k'}$ is the embedding that achieves the bounds of Theorem 3.2, and $\Phi : S \to \ell_p^{vk'}$ is defined by the direct sum $\Phi = \bigoplus_{j\in[v]} \Phi_j$, where each $\Phi_j : S \to \ell_p^k$ is given by

$$\Phi_j = \sum_{i\in I:\ i\equiv_v j} \frac{\varphi_i}{(1+\varepsilon)^{i(1-\alpha)}}.$$

The final embedding is $\Phi/M^{1/p} : S \to \ell_2^{2vk'}$ (for the same M as above), which has the required target dimension.

We need to make only small changes to the preceding proof of distortion: In the statement and proof of Lemma 4.5 and elsewhere, the dividing term $(1+\varepsilon)^{i/2}$ is replaced by $(1+\varepsilon)^{i(1-\alpha)}$. Note that the increase in value of v (and the introduction of the term $\tilde{\alpha}$), is necessary for Lemma 4.6 to hold in this setting. (In particular, the bounds on the geometric series in the proof of Lemma 4.5 (a) require, for the above choice of v, that $\tilde{\alpha} \le \alpha$, and the bounds on the geometric series in the proof of (b) require $\tilde{\alpha} \le 1-\alpha$.) No other changes to the proof are necessary, and this completes the proof of Lemma 4.4.

Comments. For l_1, the snowflake dimension can be further improved by plugging in the dimension reduction embedding of [OR02]. For $\alpha \le \frac{p}{2}$, since the α-snowflake of ℓ_p is in l_2 [DL97], one could instead embed the snowflake into l_2, reduce dimension via [JL84] or via a second snowflake from l_2 to l_2 [BRS11, GK11], and embed back into ℓ_p using [JS82]. In regards to Lemma 4.2 above, we conjecture that the target dimension in Lemma 4.4 can be reduced to $k = \tilde{O}(\varepsilon^{-3}\tilde{\alpha}^{-2}(\dim S))$ by combining randomness.

5 Clustering

Here we show that our snowflake embedding can be used to produce faster algorithms for the k-center and min-sum clustering problems. In both cases, we obtain improvements whenever $(\mathrm{ddim}/\varepsilon)^{\Theta(1)}$ is smaller than the ambient dimension.

k-center clustering. In the k-center clustering problem, the goal is to partition the input set into k clusters, where the objective function to be minimized is the maximum radius among the clusters. Agarwal and Procopiuc [AP02] considered this problem for d-dimensional set $S \subset \ell_p$, and gave an exact algorithm that runs in time $n^{O(k^{1-1/d})}$, and used this to derive a $(1+\varepsilon)$-approximation algorithm that runs in $O(nd\log k) + (k\varepsilon^{-d})^{O(k^{1-1/d})}$. Here, the cluster centers are points of the ambient space $\mathbb{R}^d$ in which S resides, chosen to minimize the maximum distance from points in S to their nearest center. The authors claim that the algorithm can be applied to the *discrete* problem as well, where all centers are chosen from S, and in fact the algorithm applies to the more general problem where the centers are chosen from a set S' satisfying $S \subset S' \subset \mathbb{R}^d$.[4] Clearly, the runtime of the algorithm can be improved if the dimension is lowered.

[4]The Euclidean core-set algorithm of [BHPI02] runs in time $k^{O(k/\varepsilon^2)} \cdot nd$, and can readily be seen to apply to all ℓ_p for constant p, $1 < p \le 2$ (see [Pan04] for a simple approach). The algorithm of [AP02] compares favorably to the core-set algorithm when d is small.

Theorem 5.1. *Given d-dimensional point set $S \subset \ell_p$ for constant p, $1 < p \leq 2$, a $(1+\varepsilon)$-approximate solution to the k-center problem on set S can be computed in time $O(nd(2^{\tilde{O}(\mathrm{ddim}(S))} + \log k)) + (k \cdot \varepsilon^{-\mathrm{ddim}(S)\log(1/\varepsilon)/\varepsilon^2})^{k^{1-(\varepsilon/\mathrm{ddim}(S))^{O(1)}}}$. This holds for the discrete case as well.*

Proof. We first consider the discrete case. Let r^* be the optimal radius. As in [AP02], we run the algorithm of Feder and Greene [FG88] in time $O(n \log k)$ to obtain a value $\tilde{r}$ satisfying $r^* \leq \tilde{r} < 2r^*$. We then build a hierarchy and extract a $\frac{\varepsilon}{2}\tilde{r}$-net $V \subset S$ in time $2^{\tilde{O}(\mathrm{ddim}(S))}nd$ (assuming the word-RAM model [BG13]). Since all points of S are contained in k balls of radius r, we have $|V| = k\varepsilon^{-O(\mathrm{ddim}(S))}$. Further, a k-clustering for V is a $(1+O(\varepsilon))$-approximate k-clustering for S. We then apply the snowflake embedding of Lemma 4.4 to embed V into $(\mathrm{ddim}(S)/\varepsilon)^{O(1)}$-dimensional ℓ_p, and run the exact algorithm of [AP02] on V in the embedded space. Since a snowflake perserves the ordering of distances, the returned solution for the embedded space is a valid solution in the origin space as well.

Turning to the general (non-discrete) case, the above approach is problematic in that the embedding makes no guarantees on embedded points not in V. To address this, we construct a set W of candidate center points in the ambient space $\mathbb{R}^d$: Recall that the problem of finding the minimum enclosing ℓ_p-ball admits a core-set of size $O(\varepsilon^{-2})$ [BHPI02]. (That is, for any discrete point set there exists a subset of size $O(\varepsilon^{-2})$ with the property that the center of the subset is also the center of a $(1+\varepsilon)$-approximation to the smallest ball covering the original set.) We take all distinct subsets $C \subset V$ of size $|C| = O(\varepsilon^{-2})$ and radius at most $\tilde{r}$, compute the center point of each subset (see [Pan04]), and add its candidate center to W. It follows that $|W| = k\varepsilon^{-O(\mathrm{ddim}(S)/\varepsilon^2)}$ and $\mathrm{ddim}(W) = \log \varepsilon^{-O(\mathrm{ddim}(S)/\varepsilon^2)} = O(\mathrm{ddim}(S)\log(1/\varepsilon)/\varepsilon^2)$. As above, we use the snowflake embedding of Lemma 4.4 to embed $V \cup W$ into $(\mathrm{ddim}(S)/\varepsilon)^{O(1)}$-dimensional ℓ_p and run the exact algorithm of [AP02] on the embedded space, covering the points of V with candidate centers from W. The returned solution is a valid solution in the origin space as well. □

Min-sum clustering. In the min-sum clustering problem, the goal is to partition the points of an input set into k clusters, where the objective function is the sum of distances between each intra-cluster point pair. Schulman [Sch00] designed algorithms for min-sum clustering under ℓ_1, ℓ_2 and ℓ_2-squared costs, and their runtimes depend exponentially on the dimension. We will to obtain faster runtimes for min-sum clustering for ℓ_0 $(1 \leq p \leq 2)$ by using our snowflake embedding as a preprocessing step to reduce dimension. Ultimately, we will embed the space into ℓ_2 using a $\frac{1}{2}$-snowflake, and then solve min-sum clustering with ℓ_2-squared costs in the embedded space; this is equivalent to solving the original ℓ_p problem. We will prove the case of ℓ_1, and the other cases are simpler.

We are given an input set $S \in \ell_1$, and set $c = 2 - \frac{1}{d'}$ for some value $d' = \left(\frac{\mathrm{ddim}(S)}{\varepsilon}\right)^{O(1)}$. We note that the $\frac{1}{c}$-snowflake of ℓ_1 is itself in ℓ_c [DL97], and also that this embedding into ℓ_c can be computed in polynomial time (with arbitrarily small distortion) using semi-definite programming [GK11], although the target dimension may be large. We compute this embedding, and then reduce dimension by invoking our snowflake embedding (Lemma 4.4) to compute a $(\frac{c}{2} = 1 - \frac{1}{2d'})$-snowflake in ℓ_c, with dimension d'. We then consider the vectors to be in ℓ_2 instead of ℓ_c, which induces a distortion of $1 + O(\varepsilon)$. Finally, we run the algorithms of Schulman on the final Euclidean space. As we have replaced the original ℓ_1 distances with their $\left(\frac{1}{c} \cdot \frac{c}{2} = \frac{1}{2}\right)$-snowflake and embedded into ℓ_2, solving min-sum clustering with ℓ_2-squared costs on the embedded space solves the original ℓ_1 problem with distortion $1 + O(\varepsilon)$.

The following lemma follows from the embedding detailed above, in conjunction with [Sch00, Propositions 14,28, full version]. For ease of presentation, we will assume that $k = O(1)$.

Lemma 5.2. *Given a set of n points $S \in \mathbb{R}^d$, a $(1 + O(\varepsilon))$-approximation to the ℓ_p min-sum k-clustering for S, for $k = O(1)$, can be computed*

- *in deterministic time $n^{O(d')}2^{2^{(O(d'))}}$.*
- *in randomized time $n^{O(1)} + 2^{2^{(O(d'))}}(\varepsilon \log n/\delta)^{O(d')}$, with probability $1 - \delta$.*

where $d' = (\mathrm{ddim}(S)/\varepsilon)^{O(1)}$.

Acknowledgements. We thank Piotr Indyk, Robi Krauthgamer and Assaf Naor for helpful conversations.

References

[ABD+08] N. Alon, M. Bădoiu, E.D. Demaine, M. Farach-Colton, M. Hajiaghayi, and A. Sidiropoulos. Ordinal embeddings of minimum relaxation: General properties, trees, and ultrametrics. *ACM Transactions on Algorithms*, 4(4), 2008.

[ABN08] I. Abraham, Y. Bartal, and O. Neiman. Embedding metric spaces in their intrinsic dimension. In *SODA '08*, pages 363–372. SIAM, 2008.

[AC06] N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. In *STOC '06*, pages 557–563. ACM, 2006.

[Ach01] D. Achlioptas. Database-friendly random projections. In *PODS '01*, pages 274–281, New York, NY, USA, 2001. ACM.

[ACNN11] A. Andoni, M. Charikar, O. Neiman, and H.L. Nguyen. Near linear lower bounds for dimension reduction in ℓ_1. In *FOCS '11*. IEEE Computer Society, 2011.

[AGMS99] Noga Alon, Phillip B. Gibbons, Yossi Matias, and Mario Szegedy. Tracking join and self-join sizes in limited storage. In *Proceedings of the Eighteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '99, pages 10–20, New York, NY, USA, 1999. ACM.

[AL11] N. Ailon and E. Liberty. An almost optimal unrestricted fast Johnson-Lindenstrauss transform. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 185–191. SIAM, 2011.

[Alo03] Noga Alon. Problems and results in extremal combinatorics, part I. *Discrete Math*, 273, 2003.

[AMS96] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 20–29, New York, NY, USA, 1996. ACM.

[And09] Alexandr Andoni. *Nearest Neighbor Search: the Old, the New, and the Impossible*. PhD thesis, MIT, 2009.

[AP02] P. K. Agarwal and C. M. Procopiuc. Exact and approximation algorithms for clustering. *Algorithmica*, 33(2):201–226, 2002.

[Ass83] P. Assouad. Plongements lipschitziens dans $\mathbf{R}^n$. *Bull. Soc. Math. France*, 111(4):429–448, 1983.

[Bal90] K. Ball. Isometric embedding in l_p-spaces. *European J. Combin.*, 11(4):305–311, 1990.

[BC05] B. Brinkman and M. Charikar. On the impossibility of dimension reduction in l_1. *J. ACM*, 52(5):766–788, 2005.

[BDH+08] M. Bădoiu, E. Demaine, M. Hajiaghayi, A. Sidiropoulos, and M. Zadimoghaddam. Ordinal embedding: Approximation algorithms and dimensionality reduction. In *APPROX '08*, pages 21–34. 2008.

[Ben62] G. Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962.

[BES06] T. Batu, F. Ergun, and C. Sahinalp. Oblivious string embeddings and edit distance approximations. In *SODA '06*, pages 792–801, 2006.

[BG13] Yair Bartal and Lee-Ad Gottlieb. A linear time approximation scheme for euclidean tsp. In *FOCS '13*, pages 698–706, Washington, DC, USA, 2013. IEEE Computer Society.

[BGK+11] Y. Bartal, L. Gottlieb, T. Kopelowitz, M. Lewenstein, and L. Roditty. Fast, precise and dynamic distance queries. In *SODA '11*, pages 840–853. SIAM, 2011.

[BHPI02] Mihai Bādoiu, Sariel Har-Peled, and Piotr Indyk. Approximate clustering via core-sets. In *STOC '02*, pages 250–257, New York, NY, USA, 2002. ACM.

[BRS11] Y. Bartal, B. Recht, and L. Schulman. Dimensionality reduction: beyond the Johnson-Lindenstrauss bound. In *SODA '11*, pages 868–887. SIAM, 2011.

[CG06] R. Cole and L. Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *38th annual ACM symposium on Theory of computing*, pages 574–583, 2006.

[CGT10] H. Chan, A. Gupta, and K. Talwar. Ultra-low-dimensional embeddings for doubling metrics. *J. ACM*, 57(4):1–26, 2010.

[DIIM04] M. Datar, N. Immorlica, P. Indyk, and V.S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 253–262. ACM, 2004.

[DKS10] A. Dasgupta, R. Kumar, and T. Sarlos. A sparse Johnson-Lindenstrauss transform. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 341–350. ACM, 2010.

[DL97] M. M. Deza and M. Laurent. *Geometry of Cuts and Metrics*. Springer-Verlag, Berlin, 1997.

[FG88] Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *STOC '88*, pages 434–444, New York, NY, USA, 1988. ACM.

[FKSV03] Joan Feigenbaum, Sampath Kannan, Martin J. Strauss, and Mahesh Viswanathan. An approximate l1-difference algorithm for massive data streams. *SIAM J. Comput.*, 32(1), January 2003.

[FS01] Jessica H. Fong and Martin Strauss. An approximate l^p difference algorithm for massive data streams. *Discrete Mathematics & Theoretical Computer Science*, 4(2):301–322, 2001.

[Gaw84] W. Gawronski. On the bell-shape of stable densities. *The Annals of Probability*, 12(1):pp. 230–242, 1984.

[GK11] L. Gottlieb and R. Krauthgamer. A nonlinear approach to dimension reduction. In *SODA '11*, pages 888–899. SIAM, 2011.

[GKL03] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, October 2003.

[Hal84] Peter Hall. On unimodality and rates of convergence for stable laws. *Journal of the London Mathematical Society. Second Series*, 30:371–384, 1984.

[HIM12] S. Har-Peled, P. Indyk, and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. Manuscript, available at `http://valis.cs.uiuc.edu/~sariel/papers/12/him/him.pdf`, 2012.

[HM06] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.

[Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[IM98] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, 1998.

[IN07] P. Indyk and A. Naor. Nearest-neighbor-preserving embeddings. *ACM Trans. Algorithms*, 3(3):31, 2007.

[Ind01] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *FOCS '01*, pages 10 – 33, 2001.

[Ind06] P. Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *J. ACM*, 53(3):307–323, 2006.

[IW05] Piotr Indyk and David Woodruff. Optimal approximations of the frequency moments of data streams. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 202–208, New York, NY, USA, 2005. ACM.

[JL84] W.B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, pages 189–206. Amer. Math. Soc., Providence, RI, 1984.

[JN09] W.B. Johnson and A. Naor. The Johnson-Lindenstrauss lemma almost characterizes Hilbert space, but not quite. In *19th Annual ACM -SIAM Symposium on Discrete Algorithms*, pages 885–891. SIAM, 2009.

[JS82] W.B. Johnson and G. Schechtman. Embedding l_p^m into l_1^n. *Acta Mathematica*, 149(1-2):71–85, 1982.

[KL04] R. Krauthgamer and J.R. Lee. Navigating nets: Simple algorithms for proximity search. In *SODA '04*, pages 791–801, January 2004.

[Kle97] J.M. Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, STOC '97, pages 599–608, New York, NY, USA, 1997. ACM.

[KOR98] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, STOC '98, pages 614–623, New York, NY, USA, 1998. ACM.

[Li08] Ping Li. Estimators and tail bounds for dimension reduction in ℓ_α $(0 < \alpha \leq 2)$ using stable random projections. In *SODA*, pages 10–19, 2008.

[LN04] J.R. Lee and A. Naor. Embedding the diamond graph in L_p and dimension reduction in L_1. *Geom. Funct. Anal.*, 14(4):745–747, 2004.

[LN05] J.R. Lee and A. Naor. Extending Lipschitz functions via random metric partitions. *Inventiones Mathematicae*, 160:59–95, 2005.

[Lug04] G. Lugosi. Concentration-of-measure inequalities. Manuscript, available at `http://www.econ.upf.edu/~lugosi/anu.ps`, 2004.

[Mat96] J. Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. *Israel J. Math.*, 93:333–344, 1996.

[MN04] M. Mendel and A. Naor. Euclidean quotients of finite metric spaces. *Advances in Mathematics*, 189(2):451 – 494, 2004.

[MS86] V. D. Milman and G. Schechtman. *Asymptotic Theory of Finite-dimensional Normed Spaces*. Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.

[Ngu13] H. Nguyen. Approximate nearest neighbor search in ℓ_p. Manuscript, available at `http://arxiv.org/pdf/1306.3601v1`, 2013.

[Nol12] J.P. Nolan. *Stable Distributions - Models for Heavy Tailed Data*. Birkhauser, Boston, 2012. In progress, Chapter 1 online at academic2.american.edu/∼jpnolan.

[OR00] R. Ostrovsky and Y. Rabani. Polynomial time approximation schemes for geometric k-clustering. In *FOCS '00*, pages 349–358. IEEE Computer Society, 2000.

[OR02] R. Ostrovsky and Y. Rabani. Polynomial-time approximation schemes for geometric min-sum median clustering. *J. ACM*, 49(2):139–156, 2002.

[Pan04] Rina Panigrahy. Minimum enclosing polytope in high dimensions. *CoRR*, 2004.

[Pan06] Rina Panigrahy. Entropy based nearest neighbor search in high dimensions. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 1186–1195, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics.

[RR07] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, 2007.

[Sch87] G. Schechtman. More on embedding subspaces of L_p in l_r^n. *Compositio Math.*, 61(2):159–169, 1987.

[Sch00] L. J. Schulman. Clustering for edge-cost minimization (extended abstract). In *32nd Annual ACM Symposium on Theory of Computing*, pages 547–555. ACM, 2000. Full version available as ECCC report TR99-035.

[Sch11] G. Schechtman. Dimension reduction in l_p, $0 < p < 2$. Manuscript, available at `http://arxiv.org/pdf/1110.2148v1.pdf`, 2011.

[SY78] K. Sato and M. Yamazato. On distribution functions of class L. *Probability Theory and Related Fields*, 43:273–308, 1978.

[Tal90] M. Talagrand. Embedding subspaces of L_1 into l_1^N. *Proc. Amer. Math. Soc.*, 108(2):363–369, 1990.

[Tal95] M. Talagrand. Embedding subspaces of L_p in l_p^N. In *Geometric aspects of functional analysis (Israel, 1992–1994)*, volume 77 of *Oper. Theory Adv. Appl.*, pages 311–325. Birkhäuser, Basel, 1995.

[VW11] S. Venkatasubramanian and Q. Wang. The Johnson-Lindenstrauss transform: An empirical study. In *ALENEX*, pages 164–173. SIAM, 2011.

[Yam78] M. Yamazato. Unimodality of infinitely divisible distribution functions of class L. *The Annals of Probability*, 6(4):pp. 523–531, 1978.

[Zol86] V.M. Zolotarev. *One-dimensional Stable Distributions*. American Mathematical Society, 1986.

A Probability theory.

The following is a simplified version of Hoeffding's inequality [Hoe63, Lug04]:

Claim A.1. *Let $X_1, \ldots, X_k$ be independent real-valued random variables, and assume $|X_i| \leq s$ with probability one. Let $\bar{X} = \frac{1}{k}\sum_{i=1}^{k} X_i$. Then for any $z > 0$,*

$$\Pr\left[|\bar{X} - \mathbb{E}[\bar{X}]| \geq z\right] \leq 2\exp\left(-\frac{2kz^2}{s^2}\right) \tag{4}$$

The following is a restatement of Bennett's inequality [Ben62, Lug04]:

Claim A.2. *Let $X_1, \ldots, X_k$ be independent real-valued random variables, and assume $|X_i| \leq s$ with probability one. Let $\bar{X} = \frac{1}{k}\sum_{i=1}^{k} X_i$ and set $\sigma^2 = \frac{1}{k}\sum_{i=1}^{k} \mathrm{Var}\{X_i\}$. Then for any $z > 0$,*

$$\Pr\left[|\bar{X} - \mathbb{E}[\bar{X}]| \geq z\right] \leq 2\exp\left(-\frac{k\sigma^2}{s^2} V\left(\frac{sz}{\sigma^2}\right)\right) \tag{5}$$

where $V(u) = (1+u)\ln(1+u) - u$ for $u \geq 0$.

Note that for $u \geq 1$, we have $V(u) = \Omega(u \log u)$, while for $0 \leq u < 1$, $V(u) = \Theta(u^2)$.